

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: The Economics of Artificial Intelligence: An Agenda

Volume Authors/Editors: Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-61333-8 (cloth); 978-0-226-61347-5 (electronic)

Volume URL: <http://www.nber.org/books/agra-1>

Conference Date: September 13–14, 2017

Publication Date: May 2019

Chapter Title: Neglected Open Questions in the Economics of Artificial Intelligence

Chapter Author(s): Tyler Cowen

Chapter URL: <http://www.nber.org/chapters/c14032>

Chapter pages in book: (p. 391 – 395)

Neglected Open Questions in the Economics of Artificial Intelligence

Tyler Cowen

Many recent writings consider artificial intelligence (AI), or more broadly “smart software,” as a transformative technology. Commonly, these writings focus on the substitution of capital for labor and the attendant domestic labor market effects. Without meaning to downplay the importance of that topic, I’d like to focus our attention on some other aspects of how artificial intelligence might affect our society.

15.1 The Distribution of Consumer Surplus

Most analyses of automation focus on the production function, but the new and cheaper outputs resulting from automation have distributional effects as well. For instance, the Industrial Revolution made food cheaper and more reliable in supply, in addition to mechanizing jobs in the factory and in the fields. A new, larger, cheaper and more diverse book market was created, and so on. Artificial intelligence, in turn, holds out the prospect of lowering prices for the outputs that can be produced by the next generation of automation. Imagine education and manufactured goods being much cheaper because we produced them using a greater dose of smart software. The upshot is that even if a robot puts you out of a job or lowers your pay, there will be some recompense on the consumer side. Internet goods such as Facebook already constitute a significant part of individuals’ time allocation, and of course they are free or very cheap at the relevant margin.

It’s worth thinking about whether the new AI-enabled outputs will be pro-

Tyler Cowen is professor of economics at George Mason University.

For acknowledgments, sources of research support, and disclosure of the author’s material financial relationships, if any, please see <http://www.nber.org/chapters/c14032.ack>.

duced at constant, increasing, or declining cost. Usually software-intensive goods tend to be produced at declining cost; namely, there is an upfront investment in the software, but at the margin additional copies are quite cheap or possibly free.

The declining cost scenario seems to have some optimistic properties. If the marginal cost is zero or near-zero, in the longer run the output price should fall considerably. In some cases, such as with social networks, the price may be zero to begin with, or perhaps negative to encourage people to join the network. Once we consider these consumption side effects, the distributional implication of an AI revolution could be more egalitarian than the job displacement effects alone would indicate.

For instance, consider the role of smartphones and cell phones in Africa today. These items have a relatively low marginal cost, and they are sold in Africa quite cheaply. They have transformed some sectors of African economies by making it much easier to manage businesses, and they also allow Africans the pleasure of communicating with each other more easily. The substitution of labor for capital in smartphone manufacturing hasn't impacted African economies much at all because Africa is not a major part of the supply chain. The more that tech production is clustered, the more that the consumption effects will be the major effects for many parts of the world.

These distribution effects may be less egalitarian if hardware rather than software is the constraint for the next generation of AI. Hardware is more likely to exhibit constant or rising costs, and that makes it more difficult for suppliers to charge lower prices to poorer buyers. You might think it is obvious that future productivity gains will come in the software area—and maybe so—but the very best smart phones, such as iPhones, also embody significant innovations in the areas of materials. A truly potent AI device might require portable hardware at significant cost. At this point we don't know, but it would be unwise to assume that future innovations will be software-intensive to the same extent that recent innovations have been.

If future AI innovations lead to very low consumer prices, this may affect our policy recommendations. Often analysts who are worried about automation call for better education and job training. Those may still be good ideas, but another approach can pay off as well. To the extent productivity is very high and prices are very low, it may suffice for workers to own some capital or natural resources. That is, wealth can serve as a substitute for income, given the extremely high purchasing power resulting from the low prices. Giving everyone some land, a birthright grant or shares in a sovereign wealth fund are options to consider, on top of whatever changes might be made to education and labor markets.

Perhaps counterintuitively, the economics of natural resources would become significantly more relevant in such a world. The scarcity of labor would matter much less, and of course robots could be used to make more

robots. You might even imagine software programs generating new products and ideas, and organizing their implementation. What would, in fact, constrain production? The answer is energy and possibly land. As scarce inputs, land and energy would determine which economies would do well and which not so well. In such a world the returns to education could be very low rather than very high.

An alternative possibility for the new scarce resource might be institutions to encourage AI-led production, such as maximally secure property rights. In that case, public choice factors would become a more significant determinant of national and regional outcomes. If “good government” is a public good of sorts, that would benefit nations and regions with especially effective norms for governance, for instance Singapore.

15.2 International Effects of an AI Revolution

Information technology also interacts with international trade. One effect of smart software is to enable more factor price equalization. It helps successful businesses become larger and also branch out internationally; for instance, it would be harder for Apple to finish off the iPhone in China if it only had the communications technologies of a few decades ago. These days, company leadership can manage an international empire by cell phone, email, and other technologies, and arguably that has led to higher investment in Chinese workers and lower investment in American and other developed country workers, especially at the lower-skilled end of the distribution.

That said, if you imagine artificial intelligence and other technologies progressing further, wage differentials might cease to be a reason to locate abroad at all. Why should the wage differential matter if the company is hardly employing any labor? As a result, there might be a reshoring of American or Western European manufacturing.

This could boost the demand for janitors here in the United States and also increase their wages, even though the number of such janitors might be small. Possibly the big income distribution effect is that artificial intelligence will be much worse for the poorer countries that can no longer industrialize through wage differentials; Dani Rodrik has labeled this phenomenon “premature deindustrialization.” At the same time, AI may be just fine for people who have the lowest wages, namely, pure manual labor jobs that can’t be outsourced at all. Information technology might be progressive at the lower end of the income distribution while hollowing out the middle, arguably a phenomenon we have seen in the United States. The biggest effects for income distribution might be across borders rather than within nations. Or, to put it another way, Africa may never have the chance to follow in the footsteps of Japan and South Korea with respect to industrialization.

From an egalitarian point of view, these distributional effects may be hard to address, precisely because they cross borders. Citizens are often willing to

support income redistribution within their nations, but they are much less likely to favor significant investments in foreign aid, especially when it is to distant nations rather than to neighbors or major trading partners.

15.3 The Political Economy of Artificial Intelligence and Income Redistribution

Discussions of artificial intelligence sometimes postulate large numbers of unemployed or underemployed people, possibly living off a guaranteed annual income or some other form of massive redistribution. On one hand, I can see the reason for considering a shift to larger cash payments. Yet the economics, politics, and sociology of guaranteed income may create problems.

If you ask which are the countries today where citizens hardly do any work, Brunei and Qatar, two resource-rich monarchies, come to mind. In each country people get a lot of money from the government, and foreign workers do much of the labor. From an analytical point of view, that is not so different from relying on robots.

The recent histories of those countries indicate that redistribution is a politically tricky concept. Imagine for instance a polity where virtually the entire gross domestic product is in some way recycled or redistributed. I expect the resulting political economy would not resemble that of Norway, as Norway without oil still would have a living standard close to that of Sweden or Denmark. Brunei or Qatar without fossil fuels likely would be *much* poorer. Given that reality, when so much of the gross domestic product (GDP) is being redistributed through politics, I wonder if this is compatible with American or Western notions of democracy. For instance, the oligarchic political forces that control the oil might make upfront offers to the interest groups that might oppose them and cement their control. Indeed those monarchies do seem to be stable, and it is far from obvious that they are evolving toward democracy. Their governments are partially benevolent toward the citizenry, but they also use a lot of the surplus to achieve their own ends, which may be religious or ideological. It seems countries that rely on fossil fuels for their GDP don't end up with the thick middle class that in the West at least partially controls the government, and is also a dominant force in our civic society and social capital. Possibly oil-rich countries do not have the economic base to sustain a version of Western-style liberal democracy, and that has something to do with so much of the GDP being recycled and redistributed. That is correlated with having a politically weak middle class and an opposition that is too easily bought off; at least that is what we observed to date in some of fossil-fuel-rich small states.

The experience of Brunei and Qatar also raises the question of what the governmental authority should be redistributing. In simple economic models, cash is redistributed to those who typically need it most. But in more comfortable settings with a lot of resource wealth, it also may be necessary

to redistribute status. That's harder to do; for the social scientist, it is also harder to model. We may need to redistribute the notion of having a meaningful job because although Qatar and Brunei have high per capita incomes, including at the median, it is not obvious to all outside observers that their citizens are happy and fulfilled.

It's possible that government "make-work" jobs will supply status to people, but there is also a danger the make-work component will be too obvious, and the resulting jobs will bring low rather than high status. In the last US presidential campaign, Hillary Clinton spoke more of redistribution and Donald Trump talked more of jobs; Trump's message seemed to be the more effective of the two.

Some desired redistributions may cross gender lines. For instance, as the population ages there will be a greater care burden for women than men, as women seem to put more time and effort into caring for their aging parents. Redistribution of money toward women may help, but at its core the problem may be one of stress rather than money per se. A change in social norms may produce a better and more effective redistribution than simply sending around checks.

If we think of caring for the elderly as a potential job with a lot of growth potential, on average women may be better at this than men, which in the labor market context serves as a penalty on being male, again to speak of the averages only. More generally, the shift toward service-sector jobs may favor women more than unskilled men. The public policies needed for many men may differ from those needed for women once again, and cash is not always the appropriate tool for recognizing those distinctions.

The general idea that in these stranger futures, what redistribution is, or has to be, is something quite different from what it is in the simple Paretian model. That is a frontier issue where we economists haven't done much work at all, but the ongoing progress of AI may make those questions all the more relevant.